

Amendments to the Specification:

Applicants hereby make amendments to the specification in order to correct typographical errors and update the patent application references with current issued patent numbers. No new subject matter has been added.

Please replace the paragraph beginning on page 2, line 20 and ending on page 3, line 3 with the following amended paragraph:

In order to generate a table index, a typical database management system must access each record and analyze the particular field or fields, i.e., key field(s) relevant to the particular index, e.g., the age field for the [[an]] age index. Using the values within the key field, the database management system orders, sorts or otherwise positions each record into a particular location as compared with the other records. This process can take considerable time because each record must be accessed from the disk and analyzed and then sorted.

Please replace the paragraph beginning on page 13, line 20 and ending on page 14, line 2 with the following amended paragraph:

Although not shown, the query manager may utilize a query optimizer, a rules engine, a cost component, a selectivity estimator and a maxdiff histogram storage area. These items are described in more detail in the United States patent application Serial No. 09/652,942 Patent No. 6, 714,938; entitled QUERY PLANNING USING A MAXDIFF HISTOGRAM, filed August 31, 2000, which is assigned to the Assignee of the present application, and which is incorporated herein by reference for all that it discloses and teaches.

Please replace the paragraph beginning on page 18, line 2 and ending on page 18, line 13 with the following amended paragraph:

With respect to the index creation system 414 shown in Fig. 4, the modules perform a number of operations or phases in order to create an index using a plurality of processing units, e.g., units 22 shown in Fig. 1. The flow 600

of logical operations performed by the index creation system 414 during index creation is shown in Fig. 6. In an embodiment, index creation manager 409 (Fig. 4) or another manager (e.g., 408 or 410) initiates flow 600 as described above. In alternative embodiments, other means are used to initiate the flow of operations shown in Fig. 6, e.g., an automatic update, addition or deletion of a record may trigger the flow shown in Fig. 6. The flow 600 relates to parallel creation of an index, i.e., using more than one processor to create the index. As such, a previous, anticipatory step may be to determine whether the system has more than one processor, or more than one available processor. Assuming this has been done and there is more than one processor available, the flow 600 may be initiated.

Please replace the paragraph beginning on page 34, line 23 and ending on page 35, line 10 with the following amended paragraph:

Once the doubly linked list is created at operation 1002, create operation 1004 creates a heap of variance values between each consecutive object in the doubly linked list. These values are then used by reduce operation 1006 to reduce the doubly linked list according to the MaxDiff Histogram Algorithm described in more detail in United States patent application Serial No. 09/652,942 Patent No. 6,714,938; entitled QUERY PLANNING USING A MAXDIFF HISTOGRAM, filed Aug 31, 2000, which is referenced above. Essentially, the MaxDiff Histogram Algorithm reduces the number of objects in the doubly linked list by determining which two objects have the lowest variance between them and then combines these two objects into a single object. Therefore, reduce operation 1006 reduces each of the various histogram lists for each of the sub-indexes to a predetermined number of objects, e.g., 200.